One single-density code, biphase Manchester, lets encoder/decoder chips serve serial data-communication systems with fewer errors, less circuitry, low overhead, and good synchronization.

Manchester II transfers data with integrity, speed

Now that Manchester coding has been embedded in single encoder/decoder chips for serial data-communication systems, the designer should understand the qualities that distinguish it from other biphase codes as well as from single-density and double-density codes in general. The differences among codes affect a media's storage capacity (bit density) and the complexity of the chip's read/write electronics. Manchester II (as it appears in such chips as the HD-6409) is distinguished by self-clocking, in which clock and data combine into a single serial-data stream (an independent track is not needed for timing). It facilitates error detection and correction; allows simpler circuitry than double-density codes and low overhead (ratio of control bits to data bits); and shows tolerance to speed variations in tape-based systems.

Code is broadly categorized as single-density or double-density. Double-density codes include delayed modulation (DM), modified-frequency modulation (MFM), group-code recording (GCR), zero modulation (ZM), enhanced nonreturn-to-zero (ENRZ), and randomized nonreturn-to-zero (RNRZ).

While the double-density codes pack twice as many bits as single-density code, they have some drawbacks. They require more complex read/write circuits, they are not self-clocking, and they require a longer preamble for synchronization, since the clock bits are not provided with each data bit. Thus, additional synchronization circuits are required. If the data consist of many short messages, the overhead for repeated preambles can be prohibitive.

Double-density codes like DM, MFM, M²FM, and GCR increase data density by replacing clock bits with data bits. MFM uses clock bits only when data bits are not present in both the preceding and current bit cell. Clocks are written at the beginning of the bit cell, and data bits are written in the middle. Since MFM is not self-clocking, the data separator is more complex than for single-density codes, but it is simpler than for most other double-density codes.

DM codes, often called phase-shift codes, come in Miller and Miller-squared (M and M²) types and are characterized by at least one transition every alternate bit cell. In Miller codes, a ONE is represented

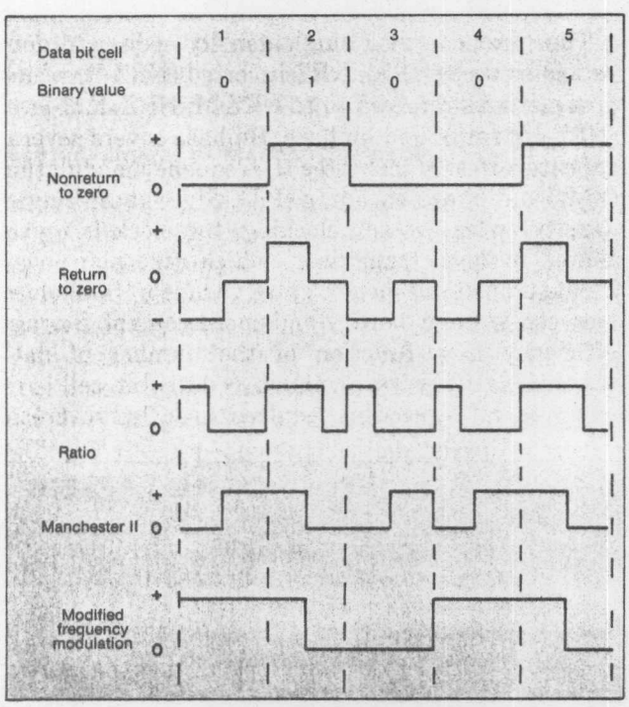

1. For a fixed binary data pattern, this timing diagram shows the various waveforms generated by five different encoding schemes.

Lester Sanders, Applications Engineer
Harris Semiconductor
P.O. Box 883, Melbourne, FL 32901

by a mid-bit-level change. A ZERO has a transition at the start of the bit cell. At least one full bit-cell period must transpire between transitions, so transitions immediately following a ZERO are delayed. In M^2 codes, a ONE transition occurs at the beginning of the bit cell, so a delay occurs whenever a ONE immediately follows a ZERO.

GCR, an outgrowth of single-density nonreturn-to-zero-inverted (NRZ-I) code, uses the same detection procedure. However, in the encode process, a 4-bit code gets converted to a 5-bit code. The 5-bit code is reconverted during the decode operation. Overall, GCR increases data storage capacity by 80% over single-density codes; however, it requires more encoder/decoder circuitry than other double-density codes.

Although they provide high densities, double-density codes also demand high-resolution record/play heads and accurate track positioning. Read/write circuits must be able to select the best write current and read gain-bandwidth. Double-density codes also have a dc component and are relatively speed-intolerant.

Single-density codes

The most common single-density codes are nonreturn to zero (NRZ); NRZ-inverted (NRZ-I), which is sometimes referred to as NRZ-M; NRZ-dual-level (NRZ-L); ratio; and biphase. Biphase covers several subcategories: Manchester II, frequency modulation (FM), and phase encoding (PE). Since these single-density codes are self-clocking, the clock is represented by level transitions, which take place even if data transitions do not. Thus, a tradeoff is involved between storage density and speed control. Storage efficiency is a function of the number of data transitions (flux reversals) per data bit cell. For example, ratio encoding requires three flux reversals per bit cell, but biphase requires two.

NRZ, return-to-zero (RZ), and biphase are categorized by the suffixes L (level), M (mark), and S (space). An -L suffix indicates that data are represented by different levels; -M and -S suffixes indicate that data are represented by the presence or absence of transitions. In codes designated -M, a ONE (defined as a mark) occurs with a level transition; ZERO is no transition. The converse is true for codes designated -S.

In NRZ codes, the logic level throughout a data cell does not change. In RZ, the logic values of data are represented by pulses that maintain their amplitude for less than an entire bit cell. In NRZ-L, a level change accompanies data changes from ONE to ZERO or ZERO to ONE. Standard RZ is an NRZ-L code in which the width of a ONE is one-half the data-bit cell. Another RZ code represents a ONE with a positive pulse during the first half of the data cell, returning to zero during the second half. It represents a ZERO with a negative pulse during the first half of the data cell, returning to zero for the second half. In NRZ codes, the information rate equals the modulation rate; however, for RZ codes, the modulation rate is double the information rate.

In NRZ-M (most often found in magnetic tape recording), a ONE by itself is represented by a full change in level. NRZ-M is often referred to as NRZ-I, where the I indicates "invert the ONEs." Another variation, known as synchronized NRZ (S-NRZ), is a modified form of NRZ-I, in which a data bit is inverted after an 8-bit data block to eliminate a possible dc component caused by too many bits in a row.

Ratio encoding, found in low-cost memory systems, is very tolerant of long-term speed variations, is self-clocking, and is relatively simple to encode. It takes three positions in the data cell for two

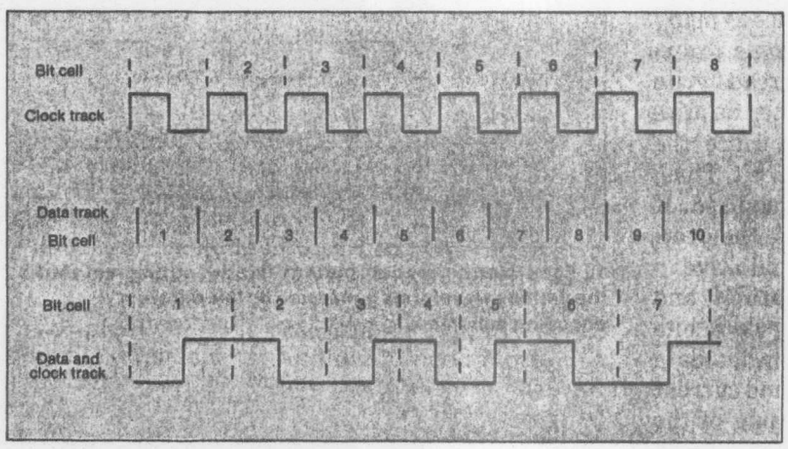

2. Comparing nonself-clocked timing (top) with self-clocked timing (bottom), these waveforms show that, over a period of time, the clock in nonself-clocked systems will appear to speed up or slow down.

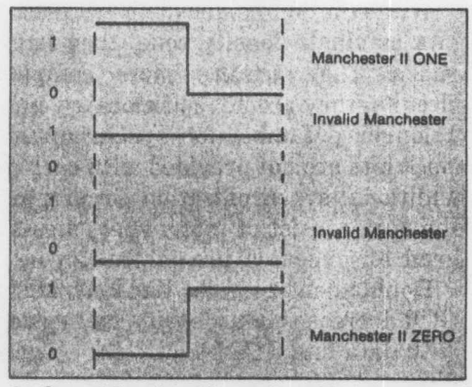

3. One of the nice things about Manchester II encoding is that errors are easily detected. Since a transition must occur in the middle of a bit cell, the only way incorrect data can be transmitted is if both the first and second half of the bit cell are inverted.

Important parameters of encoding schemes

Code	Bandwidth f_l	Bandwidth f_h	Storage efficiency	Self-clocking	Dc presence	Band-speed ratio	Tolerance to transport-speed variations: Short-term	Tolerance to transport-speed variations: Long-term	Preamble for synchronization
NRZ	0	0.5 f*	100%	No	Yes	∞			
RZ	0.25 f	1.0 f	50%	No	Yes	4	Good	Good	No
S-NRZ	0	0.5 f	80%	No	No	9	Poor	Poor	Yes
Ratio	0.75 f	1.5 f	33%	Yes	No	2	Good	Good	No
Biphase	0.5 f	1.0 f	50%	Yes	No	2	±33%	±33%	No
Double-density MFM, DM		0.5 f	100%	No	Yes	2	Poor	Poor	Yes

*Bandwidth in terms of the fundamental frequency of the data rate.

transitions. The first transition is positive, and defines the beginning of the data cell. The binary value of the cell is then determined by how long the information remains high relative to the time it is low.

Biphase encoding comes in many variants, but all are characterized by having at least one transition per bit period. In Biphase-L (Manchester II), a transition occurs at the center of the data cell, with a ONE defined as a high level for the first half of the bit period and a low for the second half of the bit period. A ZERO has the reverse levels for the bit cell. Thus, the data are defined by the direction of the transition. In Biphase-M, the transition occurs at the beginning of the data cell, with a ONE represented by a transition at mid-bit and a ZERO by the absence of the transition. Biphase-S is the exact reverse of Biphase-M.

The Table summarizes the major features for some of the popular single and double-density codes. The encoded waveforms in Fig. 1 illustrate the patterns for an identical binary input produced by each form of encoding.

Biphase stands out for simplicity

Where double-density represents overkill, biphase provides the best of all choices: simple circuitry and low overheads. Biphase codes like Manchester II—a popular code for magnetic recording and optical communication systems—are self-clocking, since the data and clock are included in a single serial data stream. For this reason, the circuitry can be simpler than in systems employing other codes. In clocked systems, the clock defines the size of the data-bit cell; however, in nonself-clocking systems, speed fluctuations cause the data track to vary relative to the speed of the clock. Over a period of time, the clock will appear to speed up or slow down and improperly define a data bit cell (Fig. 2, top). With self-clocking, since the mid-bit transition (Manchester coding has at least one transition per bit period) is the clock, everything stays synchronized (Fig. 2, bottom).

Furthermore, the mid-bit transitions of Manchester code help detect transmission errors. For example, suppose that during a bit cell, a ONE, as represented by the top row in Fig. 3, is to be transmitted. However, noise causes a logic inversion in the second half of the bit cell, so the data are transmitted as shown in the second row. Since no mid-bit transition occurs, the transmission is detected as an error by the receiving circuitry. An error would likewise be indicated if the inversion occurred in the first part of the bit cell. For an incorrect bit to be transmitted, an inversion would have to occur on both halves of the bit cell—a very unlikely event.

Take advantage of Manchester code

The HD-6409 Manchester encoder/decoder (MED), which can be programmed to operate either as a converter or a repeater, and the HD-6419, which is a fixed-repeater version of the circuit, are suitable for asynchronous serial data systems with data rates up to 1 Mbit/s.

As a converter, the 6409 independently encodes NRZ into Manchester or decodes Manchester into NRZ. In the repeater mode, the circuit accepts Manchester data streams and retransmits them with a recovered clock. Once again, the Manchester code input is decoded into NRZ.

In a low-cost digital-cassette storage system that stores about 700 kbytes of data on a standard cassette and about 200 kbytes on a minicassette, the 6409 easily realizes a data-transfer rate of 24 kbits/s, with a recording density of 800 bits/in. and an average access time of about 20 s. Design considerations for the system include the interface circuitry between the host computer and transport, the formatting of the data on the tape, and the process for encoding and decoding the data. Furthermore, error detection and correction may be provided to ensure data accuracy. All but the most elementary systems need a file structure.

The information recorded on tracks consists of

characters, files, and records. Groups of 8 to 16 bits form characters or words; two to several thousand characters constitute a record; and two to several hundred records create a typical file. In one tape format, common to digital-cassette systems (Fig. 4), the data records or blocks are stored on the blank tape, each separated by an interrecord gap (IRG).

To conform to standards for digital cassettes (ANSI X3.48-1977 or ECMA), the IRG is written with a high polarity. Each record begins with a preamble and ends with a postamble, which consists of eight bits of alternating ONEs and ZEROs, beginning with ZERO and ending with ONE. The final ONE of the postamble remains high during the IRG.

The optimum file structure depends on the nature of the information stored on the tape. Overhead efficiency is one criterion by which the file structure can be judged—that is, the number of control bits relative to the data bits. In a universal asynchronous receiver/transmitter, appending start, stop, and parity bits to seven data bits yields an efficiency of 70%. However, with the exception of its synchronization sequence, the HD-6409 MED encodes and decodes without overhead. Thus, it is transparent to any format and very flexible.

When companies do not follow the ANSI standards —reaching, instead, for a broader market—compatibility can be a problem. The methods, as well as the formats, for accessing data may differ. In one common scheme, a microcontroller searches for a desired record by counting IRGs while moving the tape at search speed. However, software can sometimes misinterpret information in the IRG as data, causing an invalid detection of a preamble. One possible solution is to put a character-recognition code immediately after a preamble and preceding the data for message verification. Software will then accept data contingent on recognition of this code. In the 6409, a 3-bit-wide synchronization sequence serves this purpose, if necessary.

Tape transports also have inertial constraints that lead to variations in tape speed, both within a machine and from machine to machine. Even high-quality transports can have speed variations of 1% or more. Errors also result if the gap in the playback head and the recorded data are not aligned at a constant angle, independently of machine and tape. The system's controller can provide four major functions in addition to establishing the link between the host computer and the transport: interface to the host; read/write amplification; encoding/decoding of the NRZ and Manchester codes; and control of the tape motion (start/stop, forward/reverse, and fast/slow operations).

Tape drives can be purchased in various degrees of completeness. A minimal drive consists of the motors, motor-drive amplifier, read/write head amplifier, and track and motion control. This last function could also include indicators and sensors for detecting the beginning-of-tape and end-of-tape,

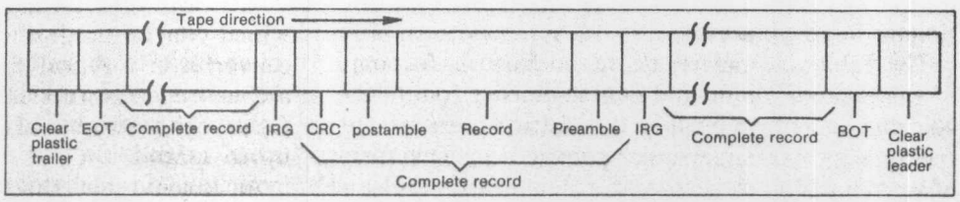

4. A typical format for magnetic recording of data requires a considerable overhead for each record—preambles, BOT header, EOT postamble, and the record itself.

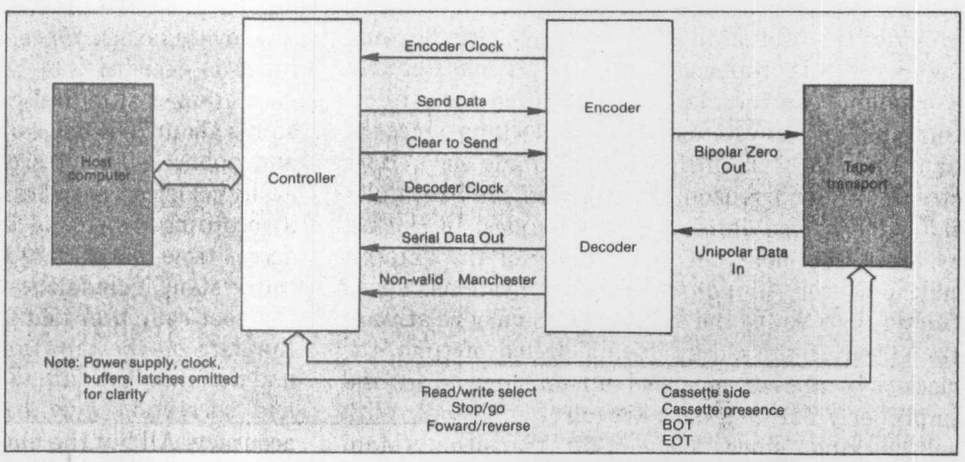

5. Using the HD-6409 Manchester encoder/decoder, this simplified data-storage system requires definition of the controller function and circuitry to interface to the host system.

Inside the Manchester encoder/decoder circuit

For high-speed yet low-power Manchester-coded communications, the HD-6409 Manchester encoder/decoder (MED) uses CMOS technology—replacing almost 250 gates. Operating at data rates from dc to 1 Mbit/s, the MED draws just 25 mW at the maximum data rate. Housed in a 20-pin DIP or 0.4-in. square leadless carrier, it has independent encoder and decoder sections, which permit full-duplex operation.

Besides the encoder and decoder circuitry, the functional sections of the chip include a free-running oscillator that is controlled by an external crystal, and the control logic (see Fig.). The control logic provides an error flag when invalid Manchester data are received, selects the clock (16 × or 32 ×), and provides encode/decode control.

For internal timing, the encoder depends on a free-running clock that operates at one or two times the data rate of the system clock. The Clear to Send (CTS) line controls the encode outputs ECLK, BOO, and BZO. ECLK is a free-running clock transmitted by the encoder to drive the external circuits that supply the NRZ data to the MED on the SD-CDS line. The complementary BOO and BZO outputs provide bipolar ONE and bipolar ZERO, respectively.

A low on the CTS line enables ECLK, BOO, and BZO; a high on the line forces BZO and BOO high and holds ECLK low. When CTS goes from high to low, a synchronization sequence is transmitted out on BOO and BZO. This sequence consists of eight Manchester ZERO bits followed by a command sync pulse (a 3-bit-wide pulse, in which the first 1.5 bits are high and the next 1.5 bits are low). A data-sync pulse is similar, but not identical (the first 1.5 bits are low and the last 1.5 bits are high); however, data-sync pulses are generated externally, not by the MED.

Serial NRZ data are shifted into the encoder section at the SD-CDS pin, upon each high-to-low transition of ECLK during the command-sync pulse. The received NRZ data are encoded into Manchester II data and transmitted out on BOO and BZO following the command-sync pulse. After the synchronization sequence, input data are encoded and transmitted out continuously, without parity checks or word framing. The Manchester data stream is sent out in inverted form, a convenience in system design since inverting line drivers can be used.

To perform the Manchester decode operation, the decoder section requires a single clock with a frequency of 16 or 32 times the desired data rate. The Speed Select (SS) pin gives the 16 × rate when low and the 32- × rate when high. A free-running clock synchronized with incoming data provides long timing-jitter margins. Manchester encoded data can be presented to the decoder in either of two forms: Bipolar ONE and ZERO inputs accept data from differential sources (such as a comparator-sensed transformer-coupled bus) or the Unipolar-Data-Input line (UDI) accepts noninverted Manchester encoded data (bipolar ZERO out of the encoder).

The decoder continuously monitors the Manchester data stream for a valid sync pattern. Although the encoder can generate only a command-sync pattern, the decoder recognizes both command and data-sync patterns. A 2-bit delay occurs between UDI, Bipolar ONE Input (BOI), or Bipolar ZERO Input (BZI), and the decoder NRZ output on the Serial Data Out (SDO) line.

Control of the decoder outputs is provided by the Reset (RST) line. When the RST line is low, SDO, DCLK, and NVM are forced low; when it is high, SDO sends out NRZ data synchronously with the recovered data clock, DCLK. The nonvalid Manchester output remains low after a low-to-high transition on RST, until a valid sync pattern is received. The DCLK signal is provided so that decoded data can be shifted into an external register on every high-to-low transition.

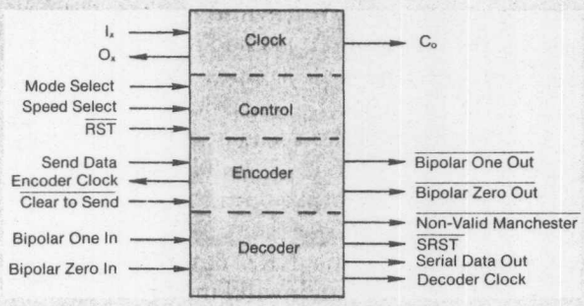

Two bit periods after an invalid Manchester bit is received on the UDI input or on the BOI and BZI inputs, the NVM line goes low synchronously with the questionable data output on the SDO line. The decoder does not re-establish proper data decoding until another sync pattern is recognized.

Additional features of the MED circuit permit it to operate in a repeater mode, under a single clock of either 16 or 32 times the desired data rate. The UDI or BOI and BZI inputs are delayed approximately one-half of a bit period and repeated as outputs BOO and BZO. The 2× ECLK is also transmitted, and is synchronous with the BOO and BZO signals. A low on CTS enables ECLK, BOO, and BZO, but does not initiate a synchronization sequence. The SD-CDS output reflects the state of the most recent sync pulse received. A high indicates a command-sync pulse, and a low indicates a data-sync pulse.

When RST is low, the outputs SDO, DCLK, and NVM are low, and SRST is set low. SRST remains low after RST goes high, and is not reset until a sync pulse and two valid Manchester bits are received with RST high. At that time, NRZ data are transmitted on the SDO line synchronously with the ×1 DCLK.

files, and the presence of cassettes. Circuitry for the interfacing, formatting, and encoding/decoding must be added along with a power supply to make the system complete. If technical resources are not available to pull the system together, complete drive systems can be purchased that simply plug into a serial or parallel port.

Figure 5 is a block diagram detailing the main signal flows in a cassette interface based on the 6409, which has a clock-recovery capability. In this system, the 6409 depends on an external clock or crystal, working at 16 or 32 times the data rate, to synchronize internal operations (see "Inside the Manchester Encoder/Decoder Circuit"). A write operation is initiated with a high-to-low transition on the Clear to Send line, which, in turn, causes the encoder to generate a synchronization sequence and transmit an Encoder Clock signal to the controller. This signal clocks the serial NRZ data from the host system into the encoder, which immediately converts the data into Manchester II code and passes it on to the cassette's write electronics.

For a read operation, the data-output line of the transport-read electronics feeds data from the tape to the unipolar data input of the MED circuit. The data are decoded and output on the Serial Data Out line. Also available is a Decoder Clock signal, which is synchronous with the NRZ data. The decoder indicates the receipt of an invalid input by signaling on the Non-Valid Manchester line. This line serves two purposes: It flags errors and masks output states that are not meant to be data (such as synchronization pulses).

To synchronize the receiving circuitry to the serial-data stream, the receiver must know where to sample the bits. Transport-speed fluctuations may make critical adjustments necessary with a time-varying data stream. Any one of several algorithms can be employed in signal-to-noise-limited channels to determine when data should be sampled. The MED uses a digital phase-locked loop (DPLL), which synchronizes at every bit period and, if necessary, adjusts for phase error between the clock and data (timing jitter). Adjustments can be made in either direction to obtain zero phase error (sampling midway between transitions).

For the MED, a free-running clock, a sample counter, an edge detector, and additional logic form the DPLL. A speed-select input chooses either a 16× or 32× data rate for the free-running oscillator. Two flags, Mid Count (MC) and End Count (EC), are initialized at bit positions 4 and 8 in the 16× mode, and positions 8 and 16 in the 32× mode (Fig. 6). The sample counter increments with the free-running oscillator, and a sample is taken when the counter reaches the EC value or a computer Final Count (FC) value. After the sample is taken, the sample counter is reset. The FC value is based on where a data transition, determined by the edge detector, occurs relative to the initialized MC and FC.

If an edge occurs either before or right when the sample counter reaches MC, the FC value is EC−1. If an edge occurs after MC but prior to ED, FC equals EC+1. If no edge is detected, a free-running sample is taken when the sample counter reaches EC.

In the synchronization sequence that precedes the data, the MED generates eight Manchester ZEROs. This procedure ensures complete synchronization of data and clock, even if the skew reaches one-quarter of a data cell. In the 32× mode, resynchronization requires eight clock cycles. Since the output is driven with the DPLL circuitry, the output-decoder clock reflects the variations in the tape speed, and is therefore used to clock data out of the decoder.□

How useful?	Circle
Immediate design application	471
Within the next year	472
Not applicable	473

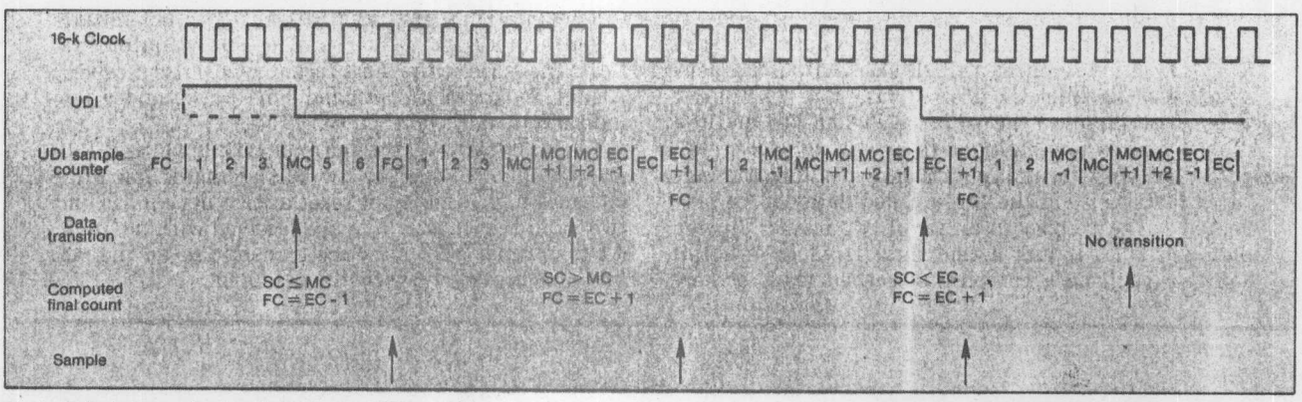

6. A digital phase-locked loop in the MED system virtually eliminates the phase error (timing jitter) caused by speed and timing variations. Shown here are synchronization examples for early, late, and no transitions for the unipolar data input in relation to the free-running counter values and the set flags.